ANALOG DEVICES

AN-282
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Fundamentals of Sampled Data Systems

A Typical DSP Sampled Data System

A block diagram of a typical sampled data DSP system is shown in Figure 3.1. Prior to the actual analog-to-digital conversion, the analog signal usually passes through some sort of signal conditioning circuitry which performs such functions as amplification, attenuation, or filtering. If the analog signal originates as a temperature, pressure, flow-rate, or force, then an appropriate sensor and transducer is required to first convert the physical quantity into an electrical voltage or current.

There are two key concepts involved in the actual analog-to digital conversion process: *discrete time sampling* and *finite amplitude resolution due to quantization.* An understanding of these concepts is vital to DSP applications.

KEY ELEMENTS OF A SAMPLED DATA SYSTEM

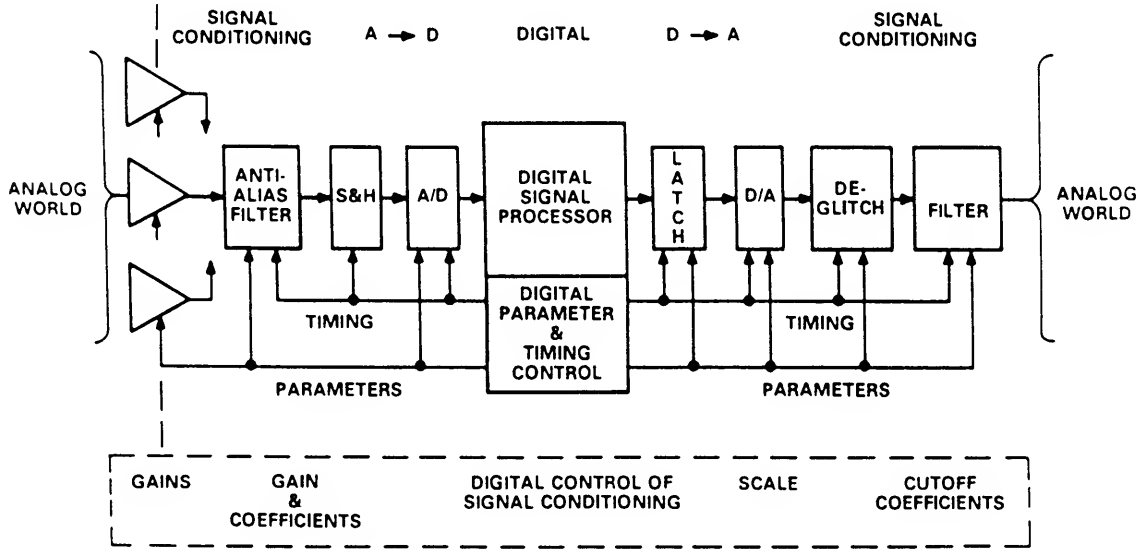

Figure 3.1

Discrete Time Sampling of Analog Signals

The concept of discrete time and amplitude sampling of an analog signal is shown in Figure 3.2. The continuous analog data must be sampled at discrete intervals, t_s, which must be carefully chosen to insure an accurate representation of the original analog signal. It is clear that the more samples taken (faster sampling rates), the more accurate the digital representation, but if fewer samples are taken (lower sampling rates), a point is reached where critical information about the signal is actually lost. This leads us to the statement of Nyquist's criteria given in Figure 3.3.

DISCRETE SAMPLING OF AN ANALOG SIGNAL

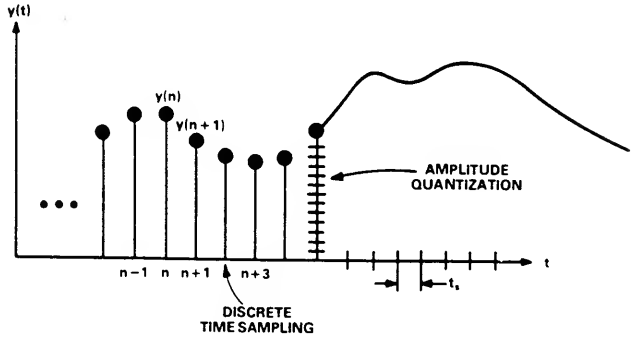

Figure 3.2

NYQUIST'S CRITERIA

- **An Analog Signal with a *Bandwidth* of f_a Must be Sampled at a Rate $f_s > 2f_a$ in Order to Avoid the Loss of Information**
- **If $f_s < 2f_a$, then a Phenomena Called Aliasing Will Occur in the Analog Signal Bandwidth**

Figure 3.3

In order to understand the implications of *aliasing* in both the time and frequency domain, first consider the four cases of a time domain representation of a sampled sinewave signal shown in Figure 3.4. In the Case 1, it is clear that an adequate number of samples have been taken to preserve the information about the sinewave. In Case 2 of the figure, only four samples per cycle are taken; still an adequate number to preserve the information. Case 3 represents the ambiguous limiting condition where $f_s = 2f_a$. If the relationship between the sampling points and the sinewave were such that the sinewave was being sampled at precisely the zero crossings (rather than at the peaks, as shown in the illustration), then all information regarding the sinewave would be lost. Case 4 of Figure 3.4 represents the situation where $f_s < 2f_a$, and the information obtained from the samples indicates a sinewave having a frequency which is lower than $f_s/2$, i.e. the out-of-band signal is *aliased* into the Nyquist bandwidth between dc and $f_s/2$. As the sampling rate is further decreased, and the analog input frequency f_a approaches the sampling frequency f_s, the aliased signal approaches dc in the frequency spectrum.

TIME DOMAIN EFFECTS OF ALIASING

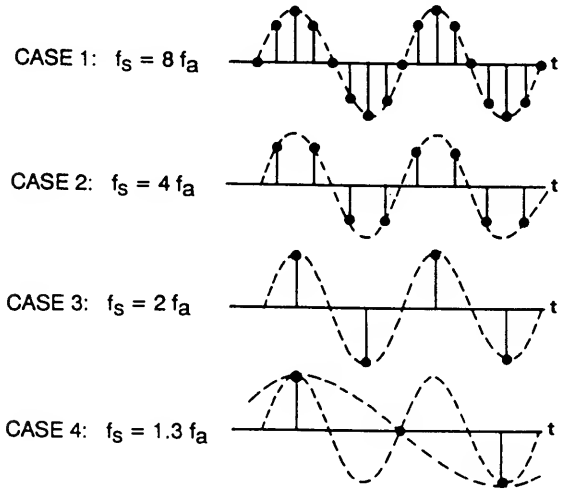

Figure 3.4

The corresponding frequency domain representation of the above scenario is shown in Figure 3.5. Note that sampling the analog signal f_a at a sampling rate f_s actually produces two alias frequency components, one at f_s+f_a, and the other at f_s-f_a. The upper alias,

FREQUENCY DOMAIN EFFECTS OF ALIASING

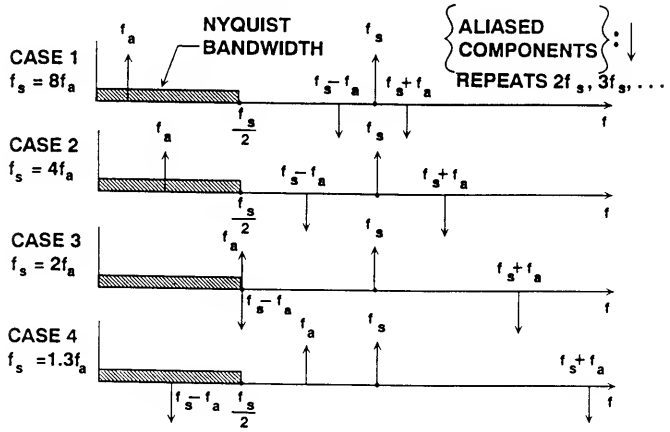

Figure 3.5

f_s+f_a, seldom presents a problem, since it lies outside the Nyquist bandwidth. It is the lower alias component, f_s-f_a, which causes problems when the input signal exceeds the Nyquist bandwidth, $f_s/2$.

It is clear from the above discussion that the ADC must be preceded by an anti-aliasing filter which has sufficient stopband attenuation at $f_s/2$ and above to prevent unwanted in-band aliasing. Aliasing may also occur from harmonics of the fundamental signal which fall outside the Nyquist bandwidth, or from unfiltered broadband noise at the ADC input.

The effects of aliasing on the dynamic range of a sampled data system are shown in Figure 3.6. The top part of the figure illustrates the desired condition at the Nyquist point, where the aliased component intersects the input signal at a point below the desired dynamic range. The lower part of the figure shows the condition where the upper-frequency dynamic range is limited by the aliased components. This condition will result in a reduction in overall signal-to-noise ratio at the higher frequencies, and could result in the distortion due to aliased out-of-band tones or harmonics as shown in Figure 3.7.

FREQUENCY DOMAIN EFFECTS OF ALIASING ON DYNAMIC RANGE

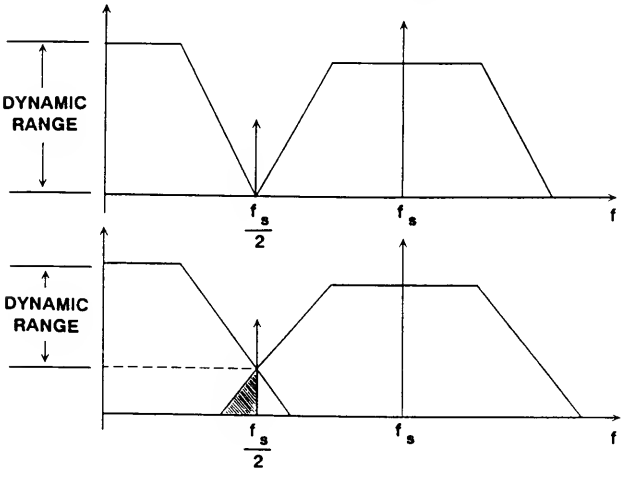

Figure 3.6

UNWANTED TONES DUE TO ALIASING

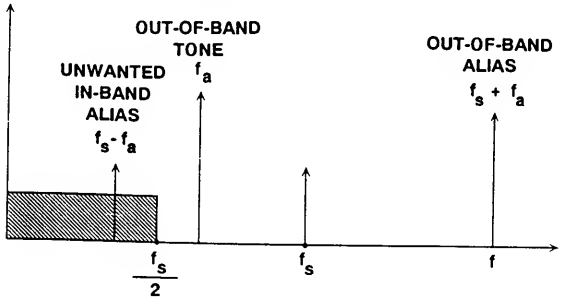

Figure 3.7

SELECTION OF ANTIALIASING FILTERS

It should be clear by now, that for a given analog input bandwidth, f_a, the requirements of the antialiasing filter are related not only to the sampling rate, f_s, but also to the desired system dynamic range. Simply stated, *dynamic range* is the ratio of the largest expected signal to the smallest signal which must be resolved, and is usually expressed in dB. At this point, we are concerned with dynamic range limitations due to aliasing. The limiting effects of ADC quantization noise and other non-linearities will be discussed shortly. The following rules of thumb will result in a filter which is somewhat overspecified, but the concepts are valid and can be refined to fit the actual system requirements.

ANTIALIASING FILTER REQUIREMENTS

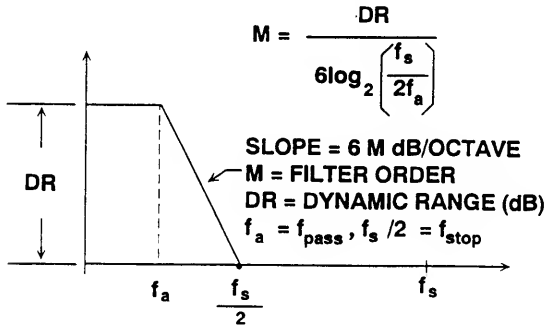

Figure 3.8

First, set the corner frequency of the antialiasing filter equal to the desired analog input bandwidth, f_a. This defines the passband of the filter, $f_{pass}=f_a$. Define the beginning of the filter's stopband, $f_{stop}=f_s/2$. Let the filter stopband attenuation be the desired upper-frequency dynamic range, DR, expressed in dB. These parameters define the transition band characteristics of the filter, i.e., it must achieve a stopband attenuation equal to the dynamic range over $\log_2(f_{stop}/f_{pass})$ octaves. The approximate order of the filter, M, (the number of poles) required to achieve this transition band slope can then be determined, since the filter rolloff is approximately 6M dB per octave. A simple example calculation is shown in Figure 3.9, where the signal bandwidth, f_a, is 3kHz, the sampling rate, f_s, is 12kHz, and a dynamic range of 60dB is required. This implies that a 10 pole filter is needed. Remember that in practice, any analog filter with more than 8 poles becomes a real design challenge, and a filter with more than 12 poles becomes almost an impossibility except

ANTIALIASING FILTER EXAMPLE

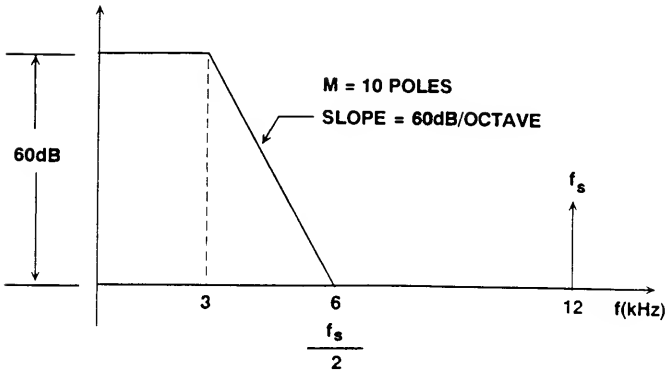

Figure 3.9

for the experienced filter designer. These considerations so far have neglected the filter's phase characteristics, and also the in-band and out-of-band ripple requirements. The addition of these parameters can make antialiasing filter design a truly formidable task.

The above rules-of-thumb for determining the complexity of the antialiasing filter assume that fullscale signals can occur at essentially all input frequencies above

Nyquist. In actual practice, this is not usually the case, and there is some natural attenuation of the signal being processed at the higher input frequencies. For instance, in the previous example, if signals at Nyquist and above were already attenuated by 12dB, then a filter stopband attenuation of only 48dB would be required at the Nyquist frequency of 6kHz. This would imply that only an 8 pole filter would be needed. This situation is illustrated in Figure 3.10.

EFFECTS OF OUT-OF-BAND ATTENUATION ON ANTIALIASING FILTER

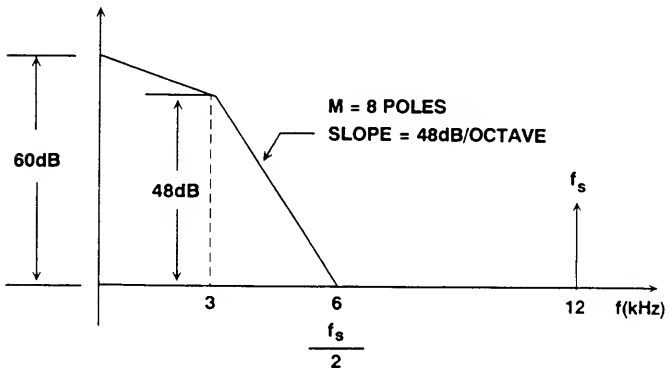

Figure 3.10

From the above discussions, it is clear that the requirements on the antialiasing filter can be relaxed at the expense of higher sampling rates (called *oversampling*). Later in the seminar, we will see that a particular class of ADCs and DACs , called *Sigma-Delta* ($\Sigma\Delta$) are inherently oversampling converters and greatly reduce the complexity of the antialiasing filter.

Oversampling and Decimation

As previously discussed, the major advantage of oversampling the input signal is the resulting simplification in the antialiasing filter requirements. Of course the downside of oversampling is that it also increases the ADC output data rate, and the DSP must be able to keep up in order to maintain real-time operation. If the data is to be transmitted in serial form, then it consequently will occupy more of the frequency spectrum. An attractive alternative makes use of both analog *and* digital filtering techniques, oversampling, and a process called *decimation*. Figure 3.11 shows the traditional case, where all the antialiasing burden lies with the analog input filter preceding the ADC. In Figure 3.12, however, the oversampling ratio, K (K is an integer), relaxes the rolloff requirement of the input analog filter by increasing the Nyquist frequency to $Kf_s/2$. The digital filter following the ADC (digital filtering will be discussed at length in Section VII) implements the antialiasing function with respect to f_s, and has sufficient stopband attenuation at $f_s/2$ to achieve the desired dynamic range. As we will see later in the seminar, digital filters having sharp cutoff characteristics with good phase response are much more easily implemented than their corresponding analog counterparts (assuming sufficient speed in the DSP). Finite Impulse Response (FIR) filters can be de-

NYQUIST SAMPLING WITH ANALOG LOWPASS FILTER

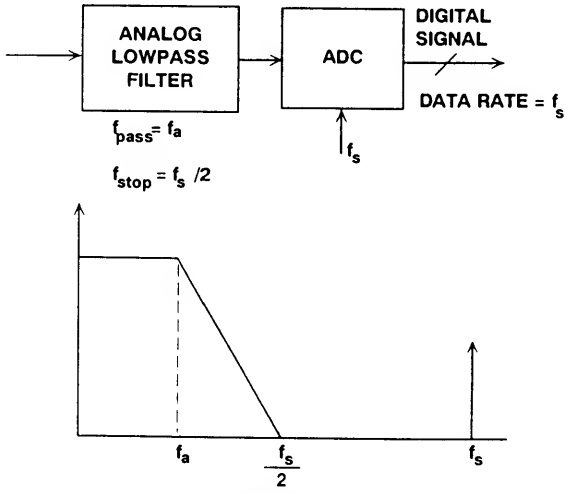

Figure 3.11

OVERSAMPLING WITH ANALOG AND DIGITAL FILTERING

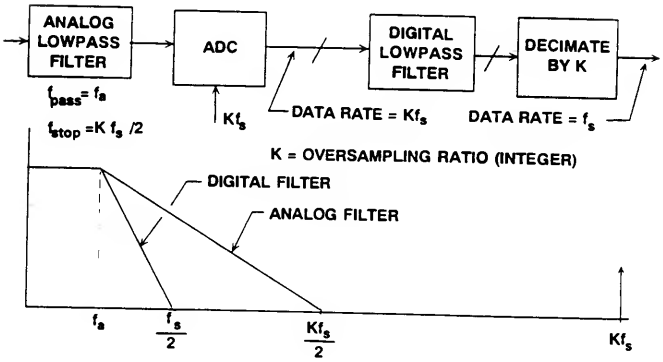

Figure 3.12

signed which have linear phase characteristics. Since the bandwidth has been reduced to $f_s/2$ by the digital antialiasing filter, the data coming out of the digital filter actually contains redundant information, and there is no need to look at every sample. In fact, it is only necessary to look at every Kth sample. This process is called *decimation*, and will be discussed in much more detail in the section of the seminar on Sigma-Delta converters (Section VI). In addition, the actual decimation can be performed by the FIR filter itself by computing a single output sample for every K input samples. This concept of oversampling and decimation is one of the most powerful concepts in real-world DSP.

Undersampling and Its Applications

In this section we will see that there are some applications in DSP where aliasing is perfectly acceptable and can be used advantageously. When the analog signal being digitized by the ADC exceeds $f_s/2$, the condition is often referred to as *super-Nyquist,* or *undersampling*. Nyquist's criteria states that the *bandwidth* (not the actual frequency) of the signal being digitized should not exceed $f_s/2$ for information to be preserved. As an example, consider a telecommunications transmultiplexer application where Frequency Division Multiplexed (FDM) data occupying the bandwidth of 60 to 108kHz is sampled at a frequency of 112kHz. Figure 3.13 shows the spectrum of the signal and the location of the aliased components. At the receiving end of the system, the filter which follows the reconstruction DAC is a bandpass rather than a lowpass and must filter out the aliased components falling between 4kHz and 52kHz as well as the component located at the sampling frequency of 112kHz.

SUPER-NYQUIST SAMPLING OF FDM SIGNAL

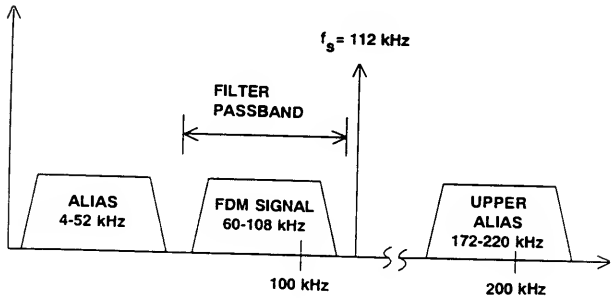

Figure 3.13

Another application for super-Nyquist operation is in the direct conversion of IF signals to baseband. Most traditional communication and radar receivers employing ADCs and DSP utilize a system in which the intermediate frequency (IF) from the front end of the receiver is down-converted or demodulated to a baseband signal by a mixer and a lowpass filter as shown in Figure 3.14. This final IF stage uses a local oscillator which is phase coherent with the signal carrier frequency. The mixer output contains a baseband signal which is proportional to the phase difference between the two inputs. Following the mixer is a lowpass filter, amplifier, and an ADC. Typical mixers have a conversion loss ranging from 4 to 6dB. In cases when the signal-to-noise ratio is limited by the front end, elimination of the mixer will improve the overall noise figure of the receiver.

ANALOG DOWNCONVERSION OR DEMODULATION

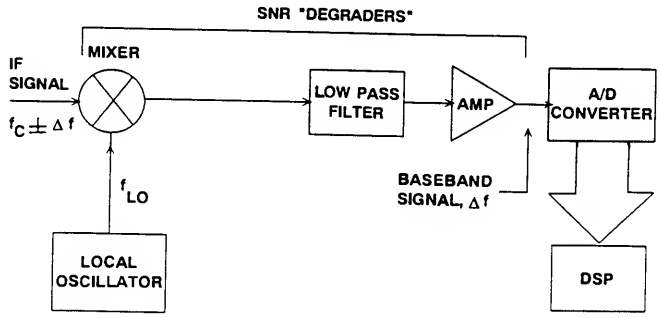

Figure 3.14

This can be accomplished (as shown in Figure 3.15) if the IF frequency is sampled at a rate equal to the local oscillator frequency. The ADC now functions as a demodulator. If the ADC samples an analog signal of the same frequency as the sampling frequency, the digitized output is a dc value. Any deviation in the analog signal from the sampling

DIRECT IF TO DIGITAL DOWN CONVERSION OR DEMODULATION

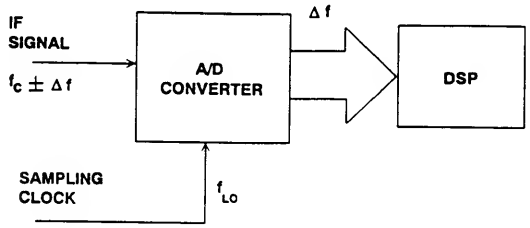

Figure 3.15

frequency looks like a *beat* frequency, Df, and the demodulation process is thereby achieved.

The data from the ADC must be processed by the DSP using an FFT which computes both the real and imaginary components of the digitized signal. This is necessary in order to preserve the phase information contained in the demodulated signal.

Operation of ADCs in a super-Nyquist environment obviously requires that the dynamic performance of the converter be known for input frequencies *above* Nyquist. The signal-to-noise ratio and harmonic distortion performance of an ADC typically degrades at higher input frequencies, so ac performance for the input frequency desired must be adequate to meet system requirements. Super-Nyquist operation typically requires an ADC which is more robust to high frequency input signals than an ADC which is specified for strictly sub-Nyquist applications.

Effects of Finite Amplitude Resolution Due to Quantization

The second major effect to be considered in a sampled data system is that of the finite amplitude resolution caused by the analog-to-digital or digital-to-analog conversion process. In this discussion, we will refer to the number of bits of the ADC (or DAC) as the converter's *resolution,* N. In the case of an ADC, the input range is divided into 2^N discrete levels, each represented by an N-bit binary word. For a DAC, the input consists of an N-bit binary word, and there are 2^N possible discrete output levels. Figure 3.16 shows the number of bits, N, the corresponding number of levels, 2^N, and the weight of the *least significant bit* (LSB) expressed as a percentage and a ratio in dB [$20 \log_{10}(2^N)$], or 6.02N dB. This ratio (whether expressed as a percentage or in dB) represents the *dynamic range* of the converter, i.e., the ratio of the largest resolvable signal to the smallest

RESOLUTION AND DYNAMIC RANGE OF ADCs AND DACs

#BITS, N	#LEVELS, 2^N	%, $100/2^N$	dB, 6N
8	256	0.4	48
10	1024	0.1	60
12	4096	0.025	72
14	16384	0.006	84
16	65536	0.0015	96
18	262144	0.0004	108
20	1048576	0.0001	120
22	4194304	0.000025	132
24	16777216	0.000006	144

Figure 3.16

resolvable signal. At this point, we should point out that the dynamic range values in Figure 3.16 represent *ideal* ADCs and DACs does not consider such ac performance limitations such as harmonic and intermodulation distortion. Neither do these values represent the theoretical signal-to-quantization noise. These topics will be discussed shortly.

Quantization Theory, Signal to Noise Ratio, and Effective Bits

The finite resolution of ADCs and DACs gives rise to a theoretical limitation to the signal-to-noise ratio (SNR) which is a function of the number of bits, N. In order to make a meaningful measurement, the ADC is stimulated with a fullscale sinewave input which is slightly below the clipping range of the converter. This gives rise to a sample-to-sample error which produces *quantization noise.* It can be shown mathematically that the rms noise voltage produced by quantization measured *within the Nyquist bandwidth* is given by the familiar expression $q/\sqrt{12}$, where q is the weight of the least significant bit (LSB) of the converter. The value for the LSB,q, can be calculated by dividing the fullscale range of the ADC or DAC by 2^N. In an ideal converter with no error sources, the theoretical rms quantization noise voltage is also independent of both the input signal amplitude and frequency. The derivation for this simple expression is given in the following reference:

W.R. Bennett, *Spectra of Quantized Signals*, BSTJ 27, pp. 446-472, July 1948

For a fullscale sinewave input, it can further be shown that the theoretical rms signal to quantization noise ratio is given by SNR = 6.02N + 1.76dB.

QUANTIZATION THEORY BASICS

- **RMS Quantization Noise in Nyquist Bandwidth, $f_s/2$:**

  $q/\sqrt{12}$

- **Fullscale Sinewave RMS Signal to RMS noise ratio in Nyquist Bandwidth:**

  $SNR = 6.02N + 1.76dB$

- **Effective Number of Bits (ENOB):**

$$ENOB = \frac{SNR_{ACTUAL} - 1.76dB}{6.02}$$

Figure 3.17

It should be noted that the rms quantization noise generally approximates broadband noise across the Nyquist bandwidth. There are certain conditions, however, where this is not true. If there is correlation between the quantization error signal and the signal being digitized, then the quantization noise may be concentrated at harmonics of the input signal rather than being spread uniformly across the bandwidth. This is most likely to occur if the input signal is a sinewave which is a subharmonic of the sampling frequency.

In testing ADCs, the SNR is usually calculated using DSP techniques while applying a pure sinewave signal to the input of the ADC as shown in Figure 3.18. The Fast Fourier Transform (FFT) processes a finite number of time samples and converts them into the frequency spectrum such as that shown in Figure 3.19 for the AD678 12-bit 200kSPS sampling ADC. The frequency spectrum is then used to calculate the SNR as well as harmonics of the fundamental input signal, very similar to an analog spectrum analyzer. The rms value of the signal is first computed. Then the rms value of all other frequency components over the Nyquist

ADC DYNAMIC TESTING

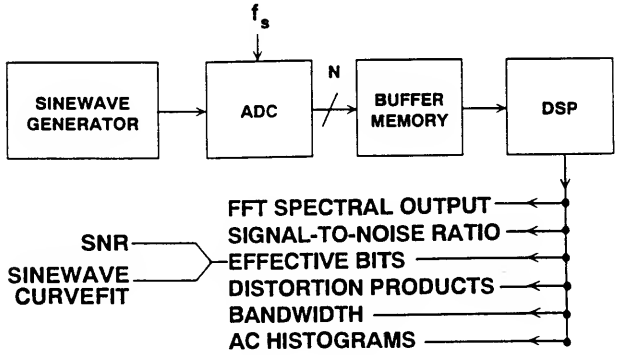

Figure 3.18

2048 POINT FFT OUTPUT FOR AD678 12-BIT, 200 kSPS ADC

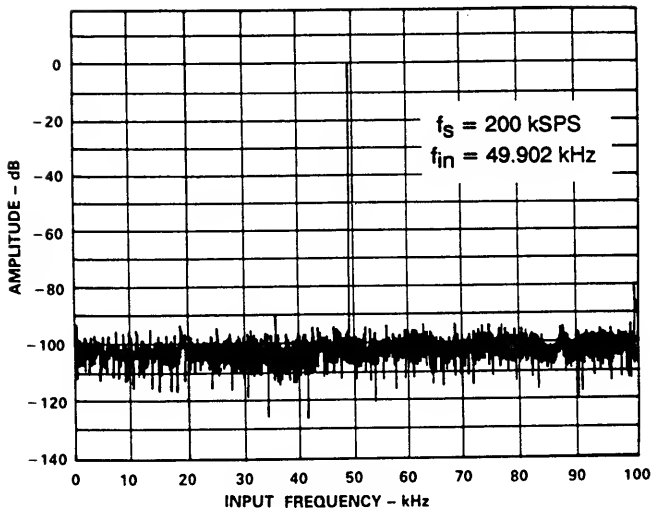

Figure 3.19

bandwidth (this includes not only noise but also distortion products) is computed. The ratio of these two quantities, expressed in dB is the SNR. Various error sources in the ADC cause the measured SNR to be less than the theoretical value, 6.02N + 1.67dB. These errors occur due to integral and differential nonlinearities, missing codes, and internal ADC noise sources. In addition, the errors typically are a function of input slewrate and therefore increase as the input frequency gets higher. In calculating the rms value of the noise, it is customary to include harmonics of the fundamental signal. This is sometimes referred to as the signal-to-noise-plus-distortion, S/(N+D), but is usually called simply SNR. A typical plot of S/(N+D) for the AD678 sampling ADC (12 bit, 200kSPS) is shown in Figure 3.20.

Another way to interpret SNR is in terms of *effective number of bits*, or *ENOBs*. The effective-bit calculation is performed by solving the SNR equation for N, given the measured value of SNR. (See Figure 3.17). For instance, a perfect 12 bit ADC would have a theoretical SNR of 74dB, corresponding to 12 effective bits. A measured SNR of 68dB, however, would correspond to 11 effective bits. This says that the performance of the actual 12 bit ADC is equivalent to that of a perfect 11 bit ADC. Figure 3.20 also shows the ENOB performance of the AD678 on the same graph as the SNR. Note that at low frequencies, the AD678 exceeds 11.4 effective bits.

Effective bits can also be measured using the *sinewave curvefit* method. In this method, a sinewave is applied to the ADC, and a number of samples are collected.

S / (N + D) AND EFFECTIVE BITS FOR AD678 12-BIT, 200 kSPS ADC

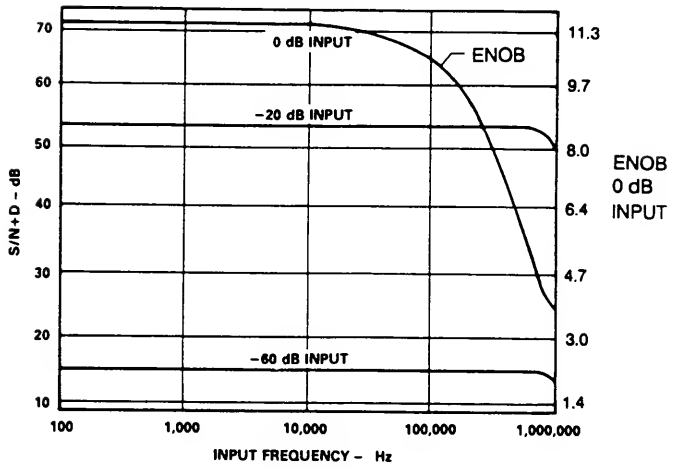

Figure 3.20

Instead of performing an FFT on the time samples, the *best-fit* sinewave to fit the data points is calculated. The sinewave amplitude, offset, frequency, and phase are chosen to minimize the rms error between the actual sinewave data points and the theoretical sinewave. Again, the theoretical rms error for a perfect ADC is $q/\sqrt{12}$. The rms error between the actual sinewave and the theoretical sinewave is computed, and the effective bits are calculated using the formula shown in Figure 3.21. The ENOB measurement using the sinewave curvefit method correlates well with that obtained using the SNR technique. If the SNR calculation is made with a signal which is less than fullscale, then a correction factor must be added as shown in order for the two methods to correlate.

CALCULATING ENOB USING SINEWAVE CURVE FITTING

- Q_A = Actual RMS Error from Fit Sinewave
- Q_T = Theoretical N-Bit RMS Error from Best Fit Sinewave

  $= q/\sqrt{12}$
- $ENOB = N - \log_2\left[\frac{Q_A}{Q_T}\right]$ Correlates to:
- $SNR = \frac{SNR_{ACTUAL} - 1.76dB + \text{Level of Signal Below FS}}{6.02}$

Figure 3.21

SELECTION OF ADC RESOLUTION BASED ON SIGNAL DYNAMIC RANGE

Selection of the proper ADC for a given application involves much more than just determining the number of bits required and the sampling rate. The dc and ac characteristics of the ADC must be examined with respect to the analog signal being processed and a proper match must be found. Inevitably, this process involves certain tradeoffs in performance and cost.

DSP APPLICATIONS AND DYNAMIC RANGE REQUIREMENTS

APPLICATION	SIGNAL BANDWIDTH	DYNAMIC RANGE	ADC # BITS
Seismology	10Hz	146dB	24
Digital Audio	20kHz	100dB	18
Echo Cancelling	4kHz	84dB	14
Speech Processing	4kHz	74dB	12
V.32 Modems	4kHz	74db	14
Ultrasound	15MHz	60dB	10
Radar	5MHz	74dB	12
Broadband Receivers	5MHz	86dB	14

Figure 3.22

Figure 3.22 shows a number of applications which are suitable for DSP processing. The approximate bandwidth and dynamic range of the corresponding signal is given. There are actually two aspects to dynamic range: *dc* and *ac*. The dynamic range corresponds to the values given in Figure 3.22 (neglecting ADC static errors). AC dynamic range, on the other hand, is related to the harmonic distortion performance of the ADC. For instance, in a digital spectral analysis application, the harmonics of a fullscale sinewave input signal limits the system's ability to resolve small signals in the presence of large signals. AC linearity is usually expressed in terms of harmonic distortion, or total harmonic distortion (THD). In a practical ADC, the number of bits may not be a good indicator of the harmonic distortion performance of the converter. AC dynamic range is less than that predicted by the SNR equation, 6.02N + 1.76dB. For these reasons, the data sheet must include both dc and ac performance specifications in order for the user to make an intelligent selection for the application.

ADC Static Transfer Characteristics

The basic specifications which describe the static performance of an ADC are given in Figure 3.23.

ADC STATIC PERFORMANCE SPECIFICATIONS

- Differential Non-Linearity (DNL)
- Integral Non-Linearity (INL)
- Missing Codes
- Gain Error
- Offset Error

Figure 3.23

In the ideal transfer function for a 3 bit ADC (Figure 3.24), the analog input signal is on the horizontal axis and the digital output is on the vertical axis. The digital output of the ADC is valid over a range of input signal. The quantum of input for a given output code is called the *width* of the code. The ideal width is exactly 1 LSB (least significant bit), but, in practice, each code-width is different from its neighbors. The deviation in the code-widths from the ideal 1 LSB value is called differential non-linearity, or DNL. A 3-bit ADC with various errors is shown in Figure 3.25. Note that the code 100 is missing because of the large DNL associated with the adjacent codes. Missing codes can produce oscillation and hunting in a closed-loop system, thus making this an important parameter to consider for ADC selection in this application.

Integral non-linearity, or INL, is usually measured with respect to the code centers. A

TRANSFER FUNCTION FOR IDEAL 3-BIT ADC

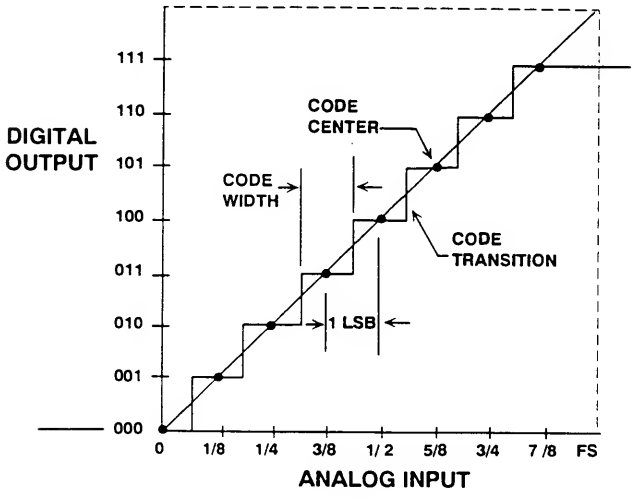

Figure 3.24

TRANSFER FUNCTION FOR NON-IDEAL 3-BIT ADC

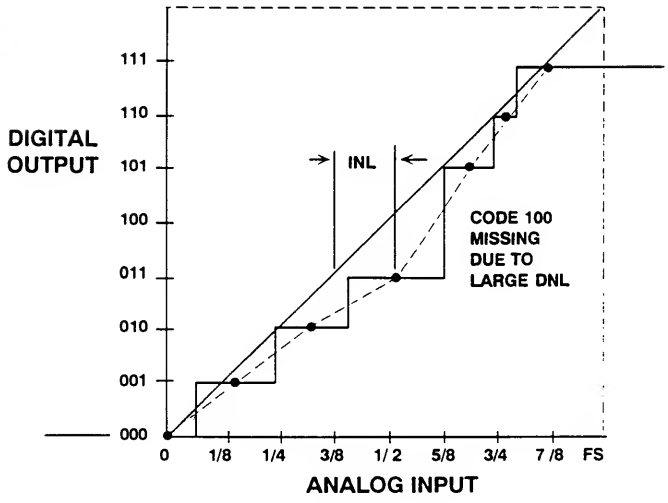

Figure 3.25

straight line is drawn through the end-points, and the worst deviation of any code center from this ideal straight line is the INL as shown in Figure 3.25. In some cases, integral non-linearity is defined with respect to a *best-fit* straight line which is typically calculated using the least-squares method.

Gain and *offset* errors apply to all codes equally and are usually trimmed out in a system using fairly traditional techniques.

DAC Static Transfer Characteristics

The basic specifications which describe the static transfer characteristics of a DAC are given in Figure 3.26.

DAC STATIC PERFORMANCE SPECIFICATIONS

- Differential Non-Linearity (DNL)
- Integral Non-Linearity (INL)
- Non-Monotonicity
- Gain Error
- Offset Error

Figure 3.26

The static transfer function for an ideal 3-bit DAC is shown in Figure 3.27. The digital input values are plotted on the horizontal axis and the corresponding analog output values on the vertical. Unlike an ADC, a DAC cannot have a missing code. There will be a discrete analog output voltage produced for each digital input code. Differential non-linearity is defined as the variation in the spacing between adjacent analog output

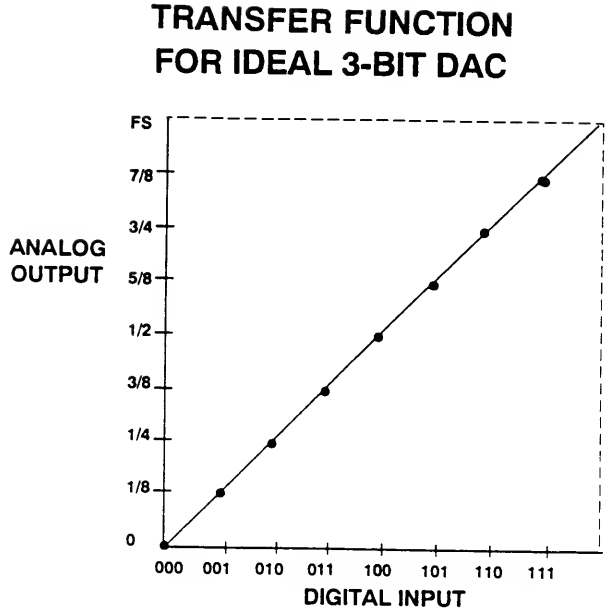

Figure 3.27

values from the ideal 1 LSB value. Excessive DNL errors can result it non-monotonic conditions as shown in Figure 3.28. A DAC is said to be non-monotonic if an increase in the digital code input causes a decrease in the analog output value. Conversely, a DAC is said to be monotonic if the slope of its transfer characteristic has the same sign over its entire range. Non-monotonic conditions can produce oscillations in a closed loop system; therefore, this specification is important in the selection of a DAC for such applications.

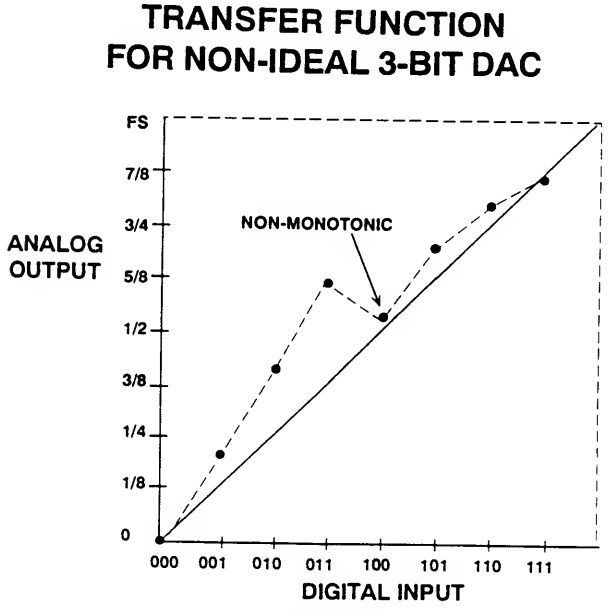

Figure 3.28

Integral non-linearity is defined as the worst case variation in any of the analog output values with respect to an ideal straight line drawn through the end points. As with an ADC, INL may also be defined with respect to a best-fit straight line.

Gain and offset definitions are similar to those for ADCs and affect each analog output value equally.

ADC Dynamic Performance

In order to be useful in most DSP applications, the ADC must have acceptable dc *and* ac performance characteristics. A listing of the most important dynamic ADC characteristics is given in Figure 3.29.

ADC
DYNAMIC SPECIFICATIONS

- **Signal-to-Noise Plus Distortion (S/N + D) Ratio and Effective Number of Bits**
- **Peak Spurious, Peak Harmonic Content, and Spurious Free Dynamic Range (SFDR)**
- **Total Harmonic Distortion (THD)**
- **Full-Power Bandwidth (FPBW)**
- **Full-Linear Bandwidth**
- **Intermodulation Distortion (IMD)**
- **Aperture Delay Time and Aperture Jitter**
- **Transient Response**
- **Overvoltage Recovery**

Figure 3.29

As we will see in a later section, there are a number of architectures which are suitable for DSP ADC designs, and most require a sample-and-hold amplifier (SHA) ahead of the actual converter as shown in Figure 3.30. Notable exceptions are flash converters and, in particular, sigma-delta converters. To fully characterize the dynamic performance of a SHA-ADC pair, they must be integrated onto the same chip, or at least offered as a complete functional unit. Otherwise, it is almost impossible to determine the overall dynamic performance of the SHA-ADC combination from the specifications on the individual devices. The requirement for complete dc and ac characterization of ADCs has led to the introduction of *sampling* ADCs which have on-board SHAs. These converters eliminate the problems of interfacing SHAs to ADCs and provide users with complete dc and ac specifications.

ADC WITH TRACK-AND-HOLD

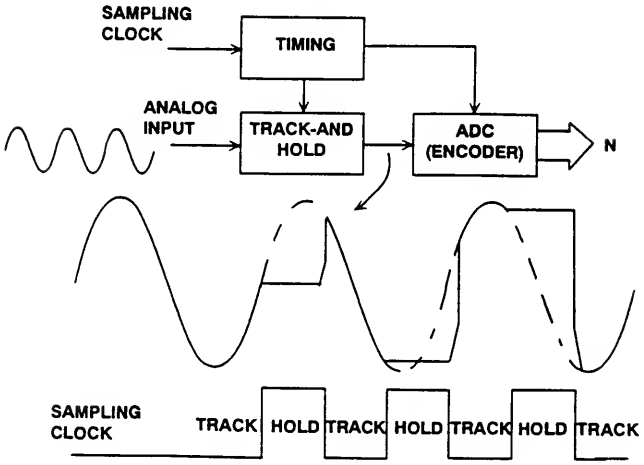

Figure 3.30

Signal-to-Noise Ratio and Effective Bits

As has been previously discussed, the signal-to-noise ratio specification is probably the most all-inclusive ac specification used in the industry today. Since it is common practice to include the effects of harmonic distortion in this measurement, S/N+D is defined as the ratio of the rms value of the measured input signal to the rms sum of all other spectral components below the Nyquist frequency, including harmonics, but excluding dc. A typical plot of S/N+D is shown in Figure 3.31 for three high speed flash ADCs. The harmonic distortion performance of the AD9617 current feedback op amp is shown on the same graph for comparison. The SNR measurement can also be expressed in effective bits, or ENOBs, as is also shown in Figure 3.31.

FLASH ADC AND OP AMP DYNAMIC PERFORMANCE

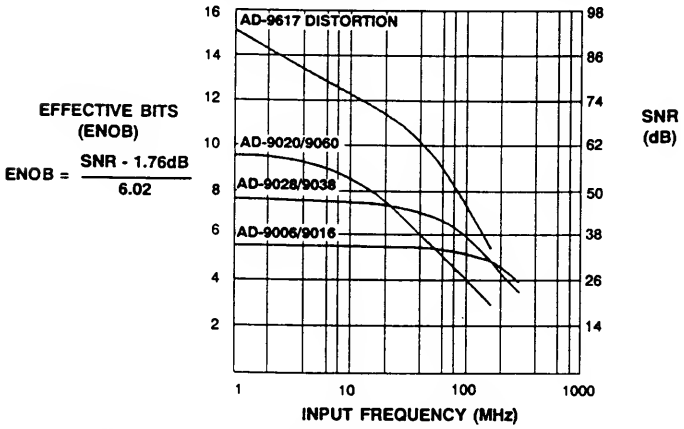

FLASH ADC	RESOLUTION	SAMPLING RATE
AD9020	10 BITS	40 MSPS
AD9060	10 BITS	60 MSPS
AD9028/9038	8 BITS	250 MSPS
AD9006/9016	6 BITS	400 MSPS

Figure 3.31

Peak Spurious, Peak Harmonic Content, and Spurious Free Dynamic Range (SFDR)

The peak spurious or peak harmonic component is the largest spectral component excluding the input signal and dc. This value is expressed in dB relative to the rms value of a fullscale input signal. The peak spurious specification is also occasionally referred to a spurious free dynamic range (SFDR). A typical plot showing the peak spurious performance for the AD678 is shown in Figure 3.32.

PEAK SPURIOUS RESPONSE FOR AD678 AT 200 kSPS, NONAVERAGED 2048 POINT FFT

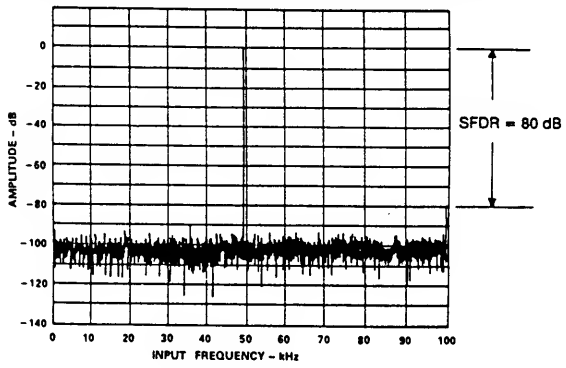

Figure 3.32

Total Harmonic Distortion (THD)

Total harmonic distortion (THD) is the ratio of the rms sum of the first six harmonic components to the rms value of a fullscale input signal and is expressed in a percentage or in dB. For input signals or harmonics that are above the Nyquist frequency, the aliased component is used. Typical THD performance for the AD678 is shown in Figure 3.33.

TOTAL HARMONIC DISTORTION, FULL-POWER BANDWIDTH, AND FULL LINEAR BANDWIDTH FOR AD678

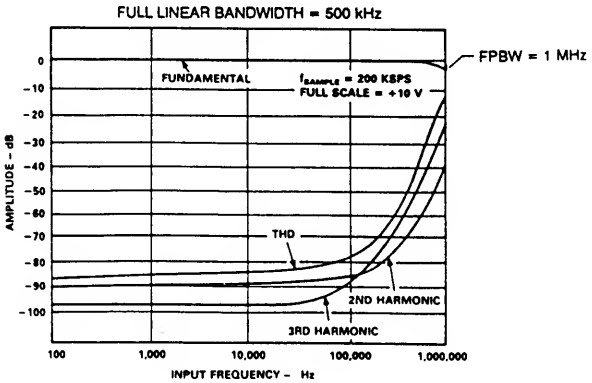

Figure 3.33

Full-Power Bandwidth

The full-power bandwidth (FPBW) of an ADC is that input frequency at which the amplitude of the reconstructed (using FFTs) *fundamental* is reduced by 3dB for a fullscale input. As can be seen from Figure 3.33, the full-power bandwidth of the AD678 is approximately 1MHz. In order to be meaningful, however, FPBW must be examined in conjunction with conjunction with SNR, ENOB, and harmonic distortion in order to determine the true dynamic performance of the ADC at the FPBW frequency.

Full-Linear Bandwidth

The full-linear bandwidth of an ADC is the input frequency at which the slewrate of the input sample-and-hold (SHA) is reached. At this point, the amplitude of the reconstructed sinewave has degraded by less than -0.1dB. Beyond this frequency, distortion of the sampled input signal increases significantly . The AD678 ADC has been designed to optimize input bandwidth, allowing it to undersample input signals significantly above the converter's Nyquist frequency. The full-linear bandwidth specification is 500kHz for the AD678 and is also shown in Figure 3.33.

Intermodulation Distortion (IMD)

Intermodulation distortion (IMD) occurs when the inputs consist of sinewaves at two frequencies, F1 and F2. Any device with nonlinearities will create distortion products, of the order (m+n), at sum and difference frequencies of mF1± nF2, where m, n = 0,1,2,3 ... Intermodulation terms are those for which m or n is not equal to zero. For example, the second order terms are (F1 + F2) and (F1 - F2), and the third order terms are (2F1 + F2), (2F1 - F2), (F1 + 2F2), and (F1 - 2F2) (see Figure 3.34). The IMD products

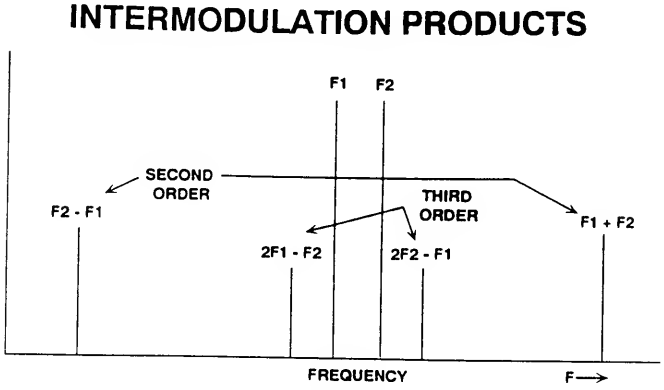

Figure 3.34

are expressed as the dB ratio of the rms sum of the measured input signals to the rms sum of the distortion terms. The two signals applied to the ADC are of equal amplitude and the peak value of their sum is -0.5dB from fullscale. The IMD products are normalized to a 0dB input signal. A typical IMD FFT plot for the AD678 is shown in Figure 3.35.

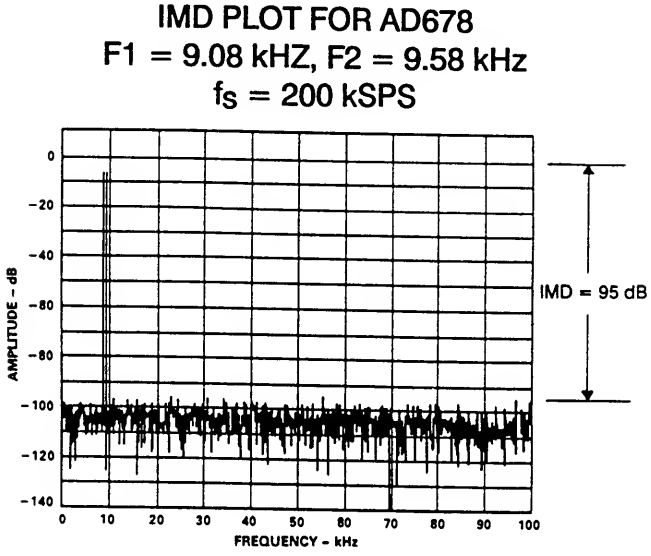

Figure 3.35

AC Linearity Plots Using Histograms

For this measurement, a fullscale sinewave is applied to the ADC, and a large number of samples are taken. The number of occurrences of each code is recorded on a histogram plot as shown in the top left-hand curve in Figure 3.36. In the case of a 12-bit converter, several million samples are required in order to achieve statistically significant results. The histogram should follow the ideal probability density distribution of a sinewave, which is shown in the top right-hand curve in Figure 3.36. The histogram data is then normalized using the sinewave probability density function to obtain the DNL plot shown in the bottom curve of the figure. Integral non-linearity can be determined by compiling a cumulative histogram. The cumulative bin widths are the transition levels. Figure 3.37 shows an ac linearity plot obtained using the histogram method for the AD7870 12-bit 100kSPS ADC digitizing a 25kHz input signal at a 100kSPS rate. The absence of large spikes in the plot shows good differential linearity. More details of the mathematics involved is given in the AD7870 data sheet .

AC LINEARITY USING HISTOGRAMS

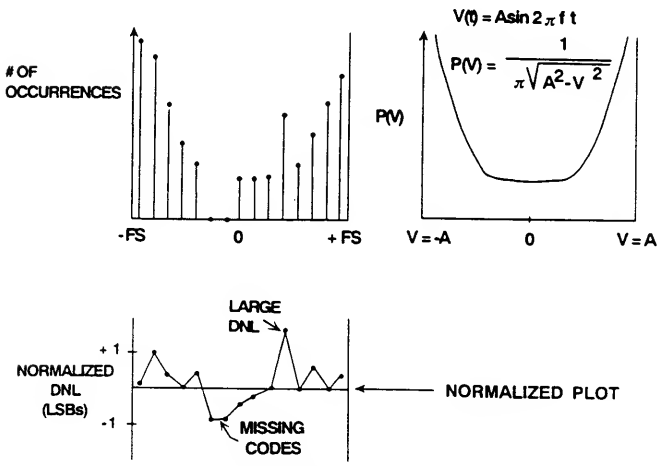

Figure 3.36

AC LINEARITY OF AD7870 12-BIT, 100kSPS ADC WITH 25 kHZ INPUT

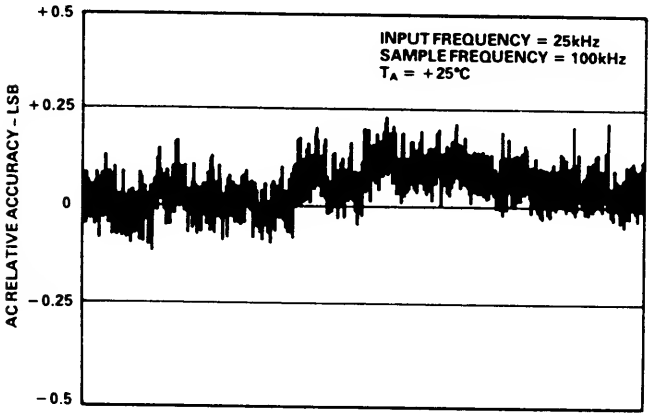

Figure 3.37

Aperture Delay Time (or Effective Aperture Delay Time)

Aperture delay time (sometimes called aperture time) is the amount of time from the leading edge of the sampling clock until the ADC actually takes the sample (see Figure 3.38). This specification is important because it helps the user to know when to apply the sampling clock with respect to the input signal timing. The variation or tolerance placed on this parameter from part to part is important in simultaneous sampling applications or other applications where the ADCs are required to track each other when processing dynamic signals.

MEASUREMENT OF EFFECTIVE APERTURE DELAY TIME

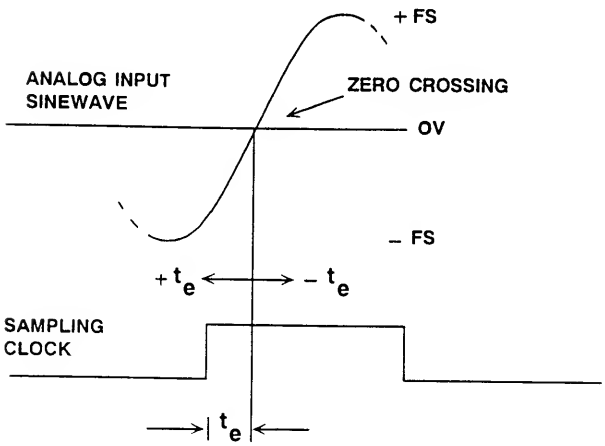

Figure 3.38

Aperture Jitter

Aperture jitter is the sample-to-sample variation in the effective point in time at which the actual sample is taken as shown in Figure 3.39. These errors generally emanate from several sources. In a practical ADC, the sampling clock is often phase-modulated by some unwanted source; the source can be wideband random noise, power line noise, or digital noise due to poor layout, bypassing, and grounding techniques. The resulting error can be expressed in terms of an rms time jitter. The corresponding rms voltage error caused by rms aperture jitter decreases the overall ADC signal-to-noise ratio. Phase jitter on the input sinewave can produce the same effect as jitter on the sampling clock.

EFFECTS OF APERTURE JITTER

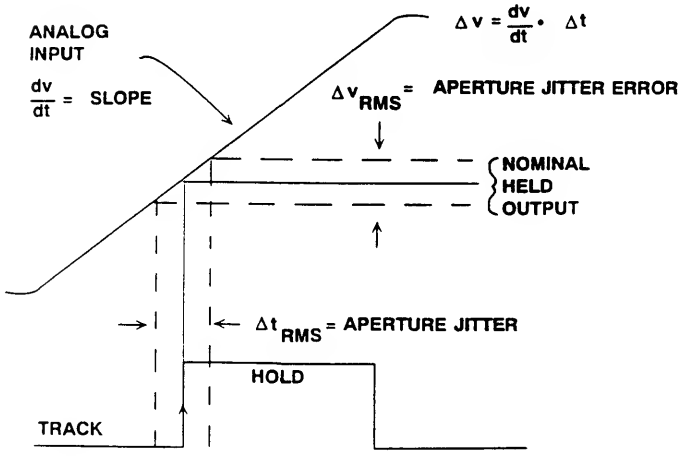

Figure 3.39

The SNR due exclusively to aperture jitter is plotted in Figure 3.40 as a function of fullscale sinewave input frequency for various values of aperture jitter. The equation for SNR due to aperture jitter is derived in Reference 1.

SIGNAL TO NOISE RATIO DUE TO APERTURE JITTER

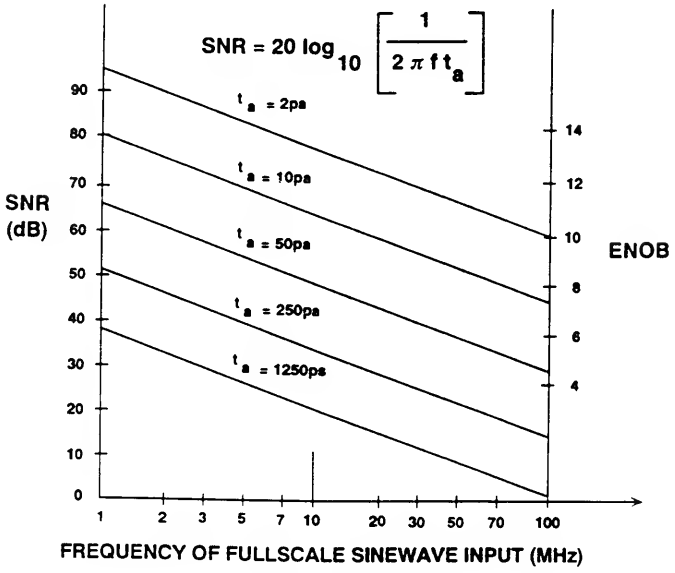

Figure 3.40

Transient Response or Settling Time

The transient response (or settling time) of an ADC is the time required for the ADC to settle to rated accuracy after the application of a fullscale step input (see Figure 3.41). This specification is critical in applications where the ADC is being driven by an analog multiplexer as shown in Figure 3.42. The multiplexer output can deliver a fullscale sample-to-sample change to the ADC input. If both the multiplexer and the ADC have not both sufficiently settled to the required accuracy, dc channel-to-channel crosstalk will result.

ADC TRANSIENT RESPONSE

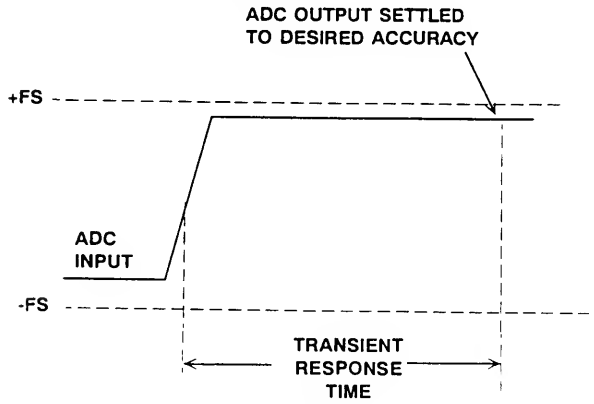

Figure 3.41

TRADITIONAL DATA ACQUISITION SYSTEM USING ANALOG MULTIPLEXER

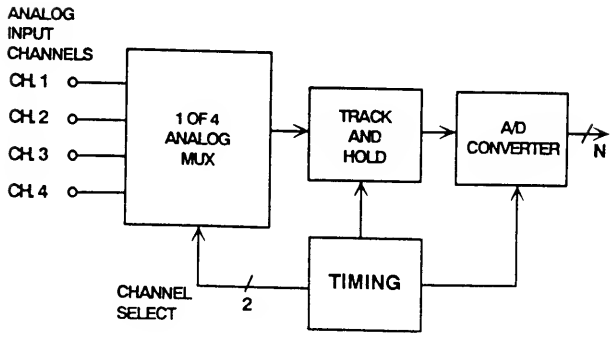

Figure 3.42

Overvoltage Recovery

Overvoltage recovery time is defined as that amount of time required for the ADC to achieve a specified accuracy, measured from the time the overvoltage signal re-enters the converter's range, as shown in Figure 3.43. This specification is usually given for a signal which is 50% outside the ADC's input range. Needless to say, the ADC should act as an ideal limiter for out-of-range signals and should produce either the positive fullscale code or the negative fullscale code during the overvoltage condition. Some converters provide over- and under-range flags to allow gain-adjustment circuits to be activated.

ADC OVERVOLTAGE RECOVERY

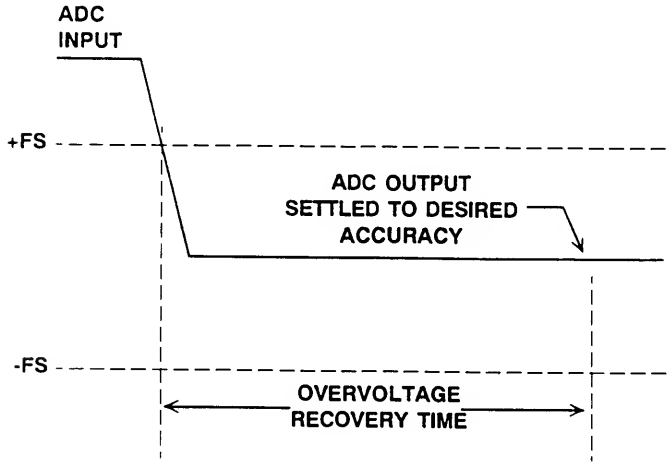

Figure 3.43

DAC Dynamic Performance

Since most DSP applications involve the eventual reconstruction of a dynamic analog signal, ac performance of DACs has become as important as ADC performance. Key DAC ac performance characteristics are given in Figure 3.45.

DAC DYNAMIC SPECIFICATIONS

- Settling Time
- Glitch Impulse Area
- Harmonic Distortion
- Signal-to-Noise Ratio
- Audio-Specific Specifications

Figure 3.44

Settling Time

Settling time of a DAC is traditionally defined as the time from the digital input transition (usually measured from the 50% point) until the DAC output settles to within a certain error band (usually 1/2 LSB) which is centered around the final value. As shown in Figure 3.45, a portion of the settling time may be due to a fixed propagation delay through the switches. If the DAC has a set of input latches or registers, the settling time should be measured from the 50% point of the latch strobe or register clock. Fullscale DAC settling time is measured for a digital input transition from 000...0 to 111...1. Midscale settling time is measured for a digital transition from 011...1 to 100...0 or 100...0 to 011...1.

It is entirely correct to define DAC settling time with respect to the output alone as shown in Figure 3.46. Settling time is measured from the time the output leaves a $\pm$1/2 LSB error band centered around the initial value until the time the output remains within a $\pm$ 1/2 LSB error band centered around the final value. The maximum DAC update rate allowable for $\pm$1/2 LSB fullscale settling time then becomes $f_{max} = 1/t_s$. Faster update rates can be used if sample-to-sample changes in the DAC input are limited to values less than fullscale.

DAC SETTLING TIME WAVEFORM

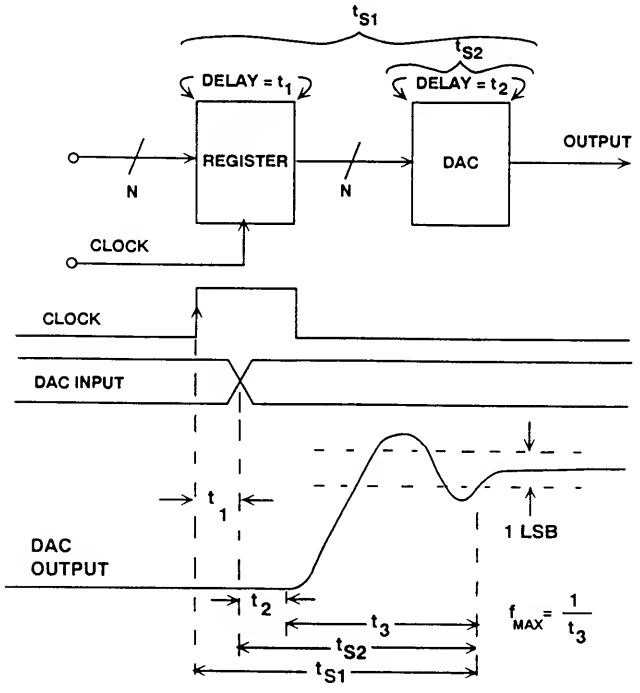

Figure 3.45

SETTLING TIME DEFINED WITH RESPECT TO DAC OUTPUT

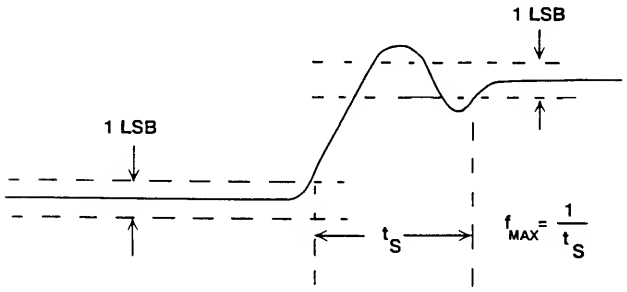

Figure 3.46

Glitch Impulse Area

Glitch impulse area is best understood by examining the waveform shown in Figure 3.47. DAC glitches occur because of digital input logic skew and unequal propagation delays through the DAC switches (a noteworthy exception to this is the sigma-delta DAC architecture to be discussed later in this seminar). The glitches are usually the largest at the midscale transition because all bits in the DAC are changing at this point. The glitch produced by the 011...1 to 100...0 transition is usually different from that produced by the 100...0 to 011...1 transition, so each must be analyzed. Glitch impulse area is simply the area of a particular glitch, and is usually measured in the units of pV-sec, therefore the fullscale output voltage of the DAC must be known in order to make meaningful comparisons between DACs. The term *glitch energy* is incorrect since the unit pV-sec is *not* a measure of energy.

From Figure 3.47 it is clear that there are six possible glitch impulse areas to deal with.

GLITCH IMPULSE WAVEFORMS

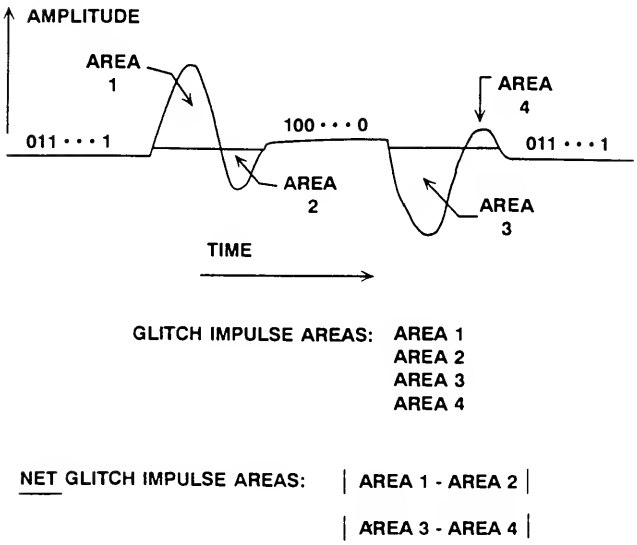

Figure 3.47

There are two glitch impulses associated with each transition. Their respective areas are designated 1,2,3, and 4. In addition, it is also useful to consider the *net glitch impulse* area associated with each of the two transitions. There are, respectively, AREA 1 - AREA 2, and AREA 3 - AREA 4. When examining the glitch impulse area specification on a DAC data sheet, it is therefore clear that there is much room for confusion unless a considerable amount of clarification is provided by the manufacturer.

Glitch impulse area remains constant regardless of filtering. Fast settling time specifications do not always imply low glitch impulse areas. The desirable situation is for the DAC to have a net glitch impulse area of zero for each of the two transitions, i.e., AREA 1 - AREA 2 = AREA 3 - AREA 4 = 0. In the ideal case, of course, each of the four areas would be zero.

HARMONIC DISTORTION

Because the net glitch impulse area is code-dependent, it will produce harmonics when the DAC is reconstructing a sinewave. A net midscale glitch occurs twice during a single cycle of the reconstructed sinewave (at each zero crossing) and, therefore, will produce a second harmonic of the sinewave as shown in Figure 3.48. Note that higher order harmonics of the sinewave which alias back into the Nyquist bandwidth are not filterable. It is difficult to predict the harmonic distortion caused by a specified net glitch impulse area, therefore, both specifications are required to adequately evaluate the dynamic performance of a reconstruction DAC.

Total harmonic distortion (THD) can be measured using DSP techniques as shown in Figure 3.49 for the AD1860 18-bit audio

EFFECTS OF DAC GLITCHES

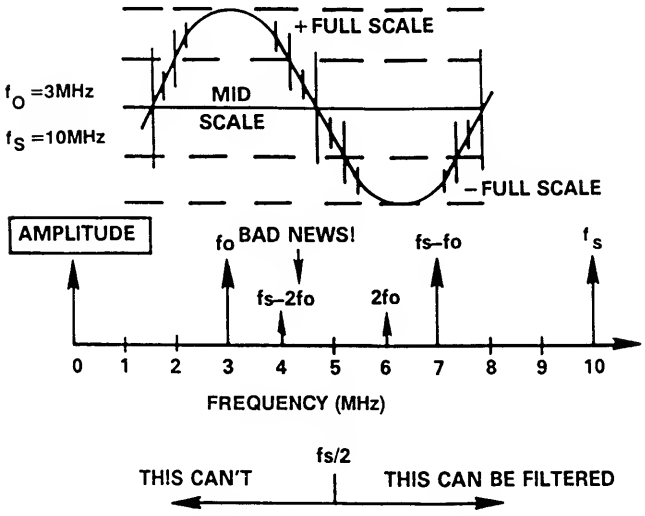

Figure 3.48

FFT TESTING OF AD1860 18-BIT AUDIO DAC

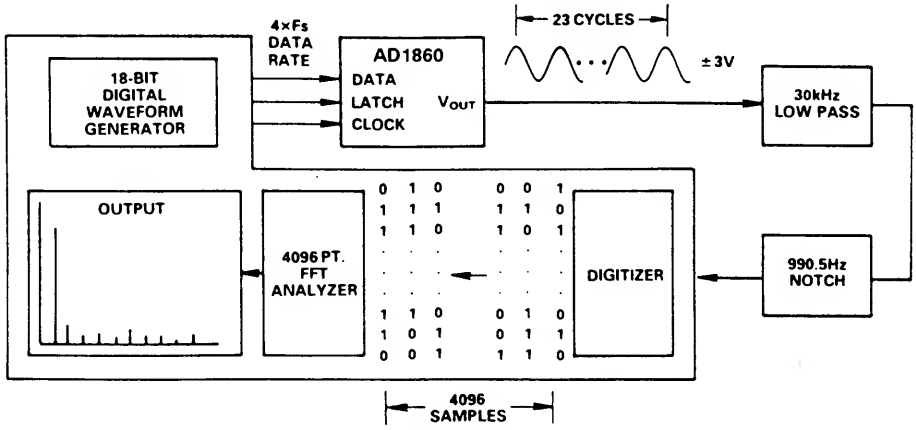

Figure 3.49

DAC. The DAC is driven with an 18-bit digital sinewave having a frequency of 990.5Hz, and the DAC update rate is 176.4kHz. The DSP digitizes 4096 samples of the output test waveform, incorporating 23 complete cycles of the sinewave. A 4096 point FFT is performed on the results of the test. The total harmonic distortion and the SNR is then calculated from the FFT results. The notch filter prevents the large-amplitude fundamental component at 990.5Hz from entering the digitizer, thereby allowing the entire digitizer range to be dedicated to processing the noise and harmonic components. Figure 3.50 shows a typical THD + noise plot for both a fullscale input and a -20dB input. It should be noted that neither a deglitcher nor an MSB trim are used in these measurements.

FFT-MEASURED
TOTAL HARMONIC DISTORTION
FOR AD1860 18-BIT AUDIO DAC

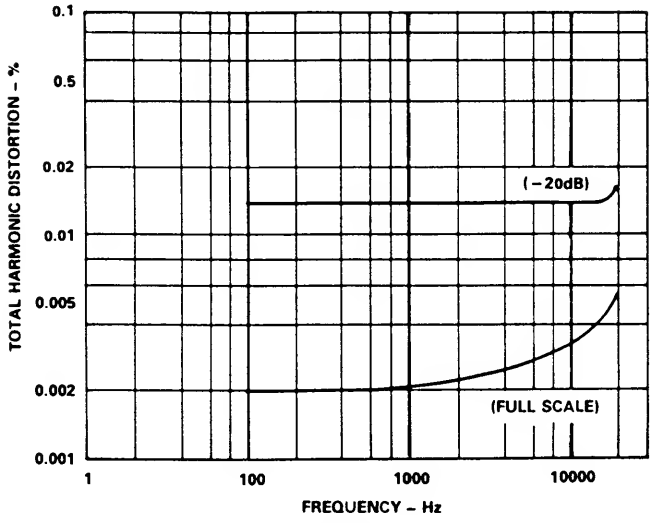

Figure 3.50

Deglitching DACs Using SHAs

SHAs can be used to deglitch DACs as shown in Figure 3.51. Just prior to latching new data into the DAC, the SHA is put into the hold mode so that the DAC switching glitches are isolated from the output. The switching transients produced by the SHA are code-independent and occur at the update frequency, hence, they are easily filterable.

SHA USED AS DEGLITCHER

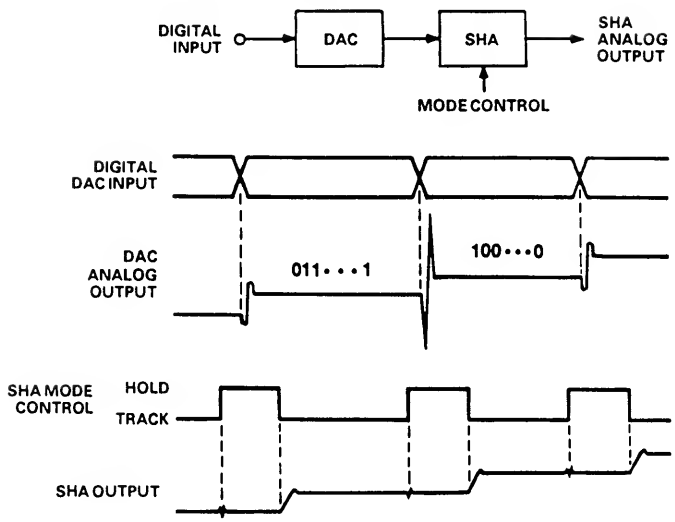

Figure 3.51

Sin(x)/x Frequency Rolloff Effect

The output of a reconstruction DAC can be visualized as a series of rectangular pulses whose width is equal to the reciprocal of the update rate as shown in Figure 3.52. Note that the reconstructed signal is down 3.92dB at the Nyquist limit with respect to the low frequency value. An inverse sin(x)/x filter is sometimes placed after the DAC to correct for this effect.

SIN X/X ROLLOFF

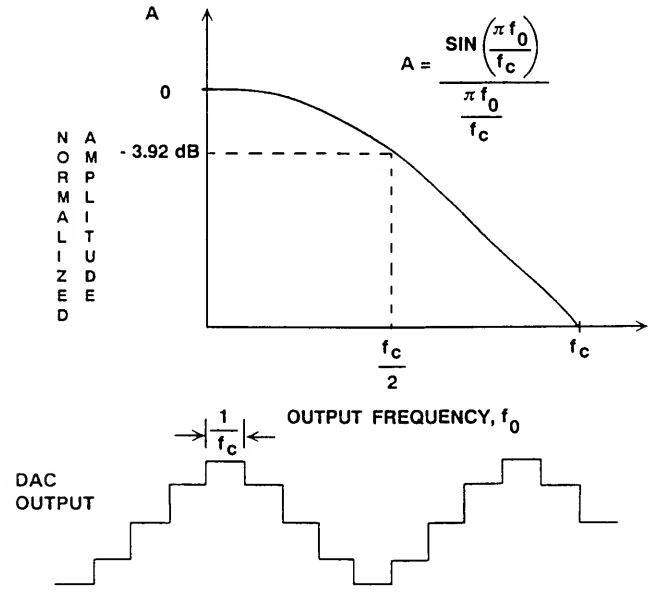

Figure 3.52

Signals were once filtered entirely in the continuous analog domain by configurations of passive components (typically inductors, resistors, and capacitors). Later active filters, with op amps for buffering and gain, provided filter designers with additional flexibility and performance, but still operated continuously on analog signals. DSP led to stable and flexible discrete-time digital filters, where sampled analog signals are processed entirely by numerical calculations with filtering algorithms—some of which cannot be realized with continuous-time analog filters.

Switched-capacitor filters (SCFs) are an intermediate class, combining both continuous- and discrete-time aspects. They are usually implemented using CMOS switches and capacitors to simulate the behavior of resistors, therefore, many filter architectures can be realized entirely by a monolithic device without the need for external components. SCFs are particularly useful for voice and audio bandwidth signal applications in conjunction with DSP technology. Since SCFs are *sampling* devices, all the concepts of discrete time sampling apply to their use: Nyquist's theorem, aliasing, etc.

FILTERING TECHNIQUES

- **Crystals, SAWs**
- **Passive Components (R, L, and C)**
- **Active Filters (R, C, and Op-Amps)**
- **Switched Capacitor Filters (CMOS Switches and Capacitors Replace Resistors)**
- **Digital Filters (Numerical Realizations Which May Have No Analog Counterpart)**

Figure 3.53

The fundamental concept of a switched capacitor acting as a resistor can best be understood from the concept of charge transfer as shown in Figure 3.54. If the capacitor is switched from V_1 to V_2, an instantaneous charge transfer occurs, $\Delta Q = C(V_1 - V_2)$, either into or out of V_2. This assumes that C has no series resistance, and that V_1 and V_2 are ideal voltage sources. If the switch is thrown back and forth at a clock frequency, f_s (having a period T), then an average current, i, flows between V_1 and V_2 having a value $i = \Delta Q/T = C\Delta V/T$. The equivalent resistance, "R", that would give the same average current is given by :

$$\text{"R"} = \Delta V/i = T/C = 1/(Cf_s).$$

SWITCHED CAPACITOR "RESISTOR"

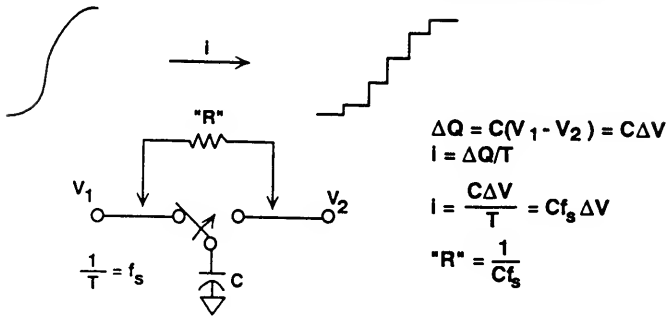

Figure 3.54

In an integrated circuit, the single-pole double-throw switch is implemented using CMOS switches driven by a non-overlapping two-phase clock as shown in Figure 3.55. A requirement for this technique to work is that the switches have very low on resistance and very high off resistance. This is precisely what CMOS technology offers.

CMOS IMPLEMENTATION OF SWITCHED CAPACITOR

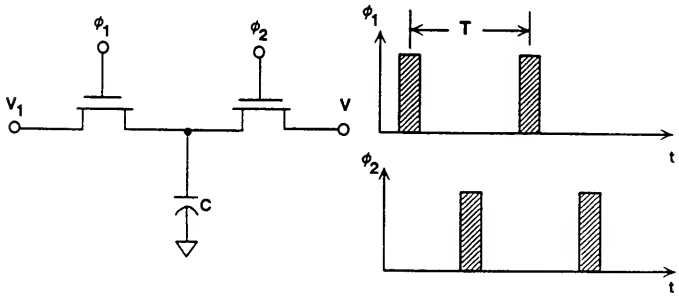

Figure 3.55

Using this SC resistor-equivalent, many conventional passive and active filter configurations can be realized. Figure 3.56 shows a single-pole passive RC filter and its SCF equivalent. The -3dB frequency of the RC filter is $1/(2\pi R_1 C_1)$. For the SCF version,

$$f_{3dB} = f_s C_1/(2\pi C_2).$$

SC EQUIVALENT OF PASSIVE RC NETWORK

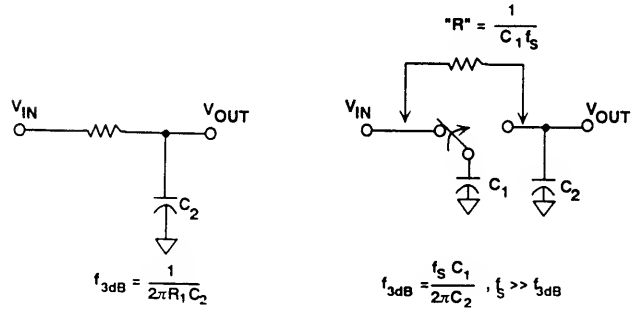

Figure 3.56

Note that for the SCF version, the *bandwidth depends on the sampling rate and the ratio of the capacitor values.* A major assumption which must be made is that $f_s >> f_{3dB}$ (typically 50 to 100) to minimize the effects of time-sampling and charge-sharing. Using the SCF concept, critical frequencies are therefore determined by capacitor ratios and the sampling clock frequency, both of which can be made precise and drift free.

Audio and voiceband filtering with SC filters can greatly reduce passive component physical size. To implement audio filters, a resistance on the order of 10MΩ is required if a monolithic capacitor of reasonable size (~ 10pF) is to be used. This value of resistance is easily achieved by switching a 1pF capacitor at a 100kHz rate, requiring a silicon area of approximately 0.01mm². If the 10MΩ resistor were implemented using polysilicon or diffusion, the area required would be at least 100 times larger.

SWITCHED CAPACITOR FILTER ADVANTAGES

- **Filter Bandwidths Proportional to Capacitance Ratios Not Absolute Values**
- **Filter Bandwidths Variable with Clock Frequency**
- **Defined Like Classic Analog Filters**
- **Low Values of Capacitance Required for Audio Frequencies: 1pF Capacitor Switched at 100kSPS = 10 MΩ "Resistance"**
- **SCFs Ideally Suited to DSP CMOS Processes**

Figure 3.57

By using SC resistors in conjunction with other capacitors and op amps, it is possible to realize many of the circuit configurations used in conventional RC active filters. Unlike digital filters, SC filters may be defined exactly like analog filters. A first-order continuous-time active lowpass RC filter and its SC counterpart are shown in Figure 3.58.

Since they sample analog signals, SC filters must usually be preceded by a continuous-time antialiasing prefilter to eliminate spectral components above the Nyquist frequency. Since the SC filter sampling rate is usually much higher than its passband, a single or double pole RC filter is usually sufficient for this purpose.

FIRST ORDER ACTIVE LOWPASS RC FILTER AND SCF EQUIVALENT

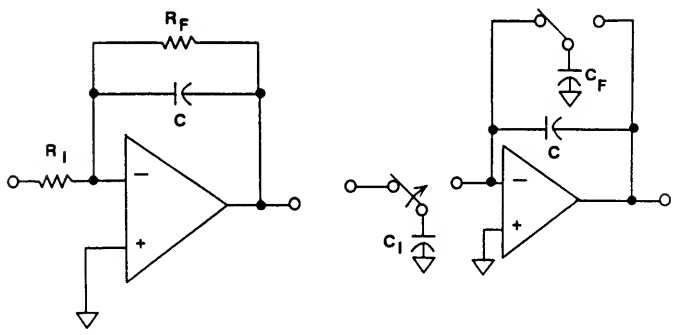

Figure 3.58

Differential amplifiers are often used in analog circuits to achieve good common mode rejection of unwanted signals such as power line noise, etc. The same principles can be used in designing switched capacitor filters. Figure 3.59 shows an active differential integrator and its switched capacitor equivalent. In addition to providing good CMRR to noise, the differential configuration also provides common mode rejection to the transients caused by the operation of the switches. Switched capacitor integrators are often used in the modulator circuits of a Sigma-Delta ADCs as will be discussed later in this seminar.

ACTIVE DIFFERENTIAL INTEGRATOR AND SCF EQUIVALENT

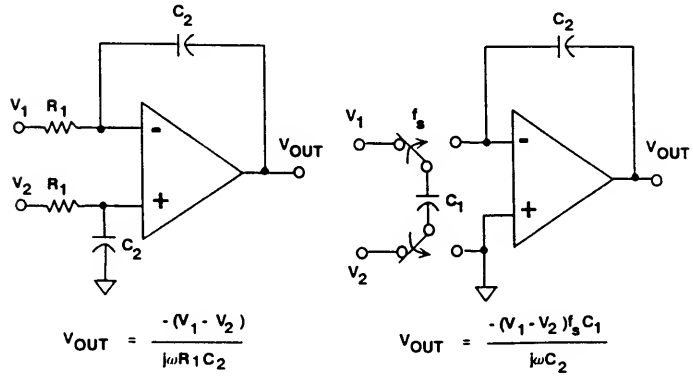

Figure 3.59

Switched capacitor filters are subject to several limitations and error sources. Their usefulness is limited to frequencies in the audio bandwidth, since sampling rates greater than a few hundred kilohertz cannot be readily achieved with current CMOS technology. The switched capacitors and op amps introduce random noise, and leakage currents can produce offset errors. Clock feedthrough from the switches themselves can produce synchronous errors. Finally, since SCFs are sampling devices, large oversampling ratios are usually required in order to prevent errors due to aliasing.

SWITCHED CAPACITOR LIMITATIONS AND ERROR SOURCES

- **Limited to Lower Frequencies**
- **Noise, Offset, and Distortion**
- **Clock Feedthrough from Switches**
- **Must Obey the Laws of Nyquist (Requires Antialiasing Filter)**

Figure 3.60

REFERENCES

1. Frederic J. Harris, "On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform", **IEEE Proceedings**, Vol. 66, No. 1, Jan. 1978, pp. 51-83.

2. Joey Doernberg, Hae-Seung Lee, David A. Hodges, "Full Speed Testing of A/D Converters", **IEEE Journal of Solid State Circuits**, Vol. SC-19, No. 6, Dec. 1984, pp. 820-827.

3. James R. Andrews, Barry A. Bell, Norris S. Nahman, and Eugene E. Baldwin, "Reference Waveform Flat Pulse Generator", **IEEE Transactions on Instrumentation and Measurement**, Vol. IM-32, No. 1, March 1983, pp. 27-32.

4. Brendan Coleman, Pat Meehan, John Reidy and Pat Weeks, "Coherent Sampling Helps When Specifying DSP A/D Converters", **EDN**, October 15, 1987, pp. 145-152.

5. Howard K. Schoenwetter, "A Programmable Voltage Step Generator for Testing Waveform Recorders", **IEEE Transactions on Instrumentation and Measurement**, Vol. IM-33, No. 3, Sept. 1984, pp. 196-200.

6. Robert W. Ramirez, **The FFT: Fundamental and Concepts**, Prentice-Hall, 1985.

7. R.B. Blackman and J.W. Tukey, **The Measurement of Power Spectra**, Dover Publications, New York, 1958.

8. James J. Colotti, "Digital Dynamic Analysis of A/D Conversion Systems Through Evaluation Software Based on FFT/DFT Analysis", **RF Expo East 1987 Proceedings**, Cardiff Publishing Co., pp. 245-272.

9. Sid Kaufman, "Multistage Error Correcting A/D Converters", **Electronic Products**, April 18, 1983, pp. 103-110.

10. **HP Journal**, Nov. 1982, Vol. 33, No. 11

11. HP Product Note 5180A-2.

12. **HP Journal**, April 1988, Vol. 39, No. 2.

13. **HP Journal**, June 1988, Vol. 39, No. 3.

14. Dan Sheingold, Editor, **Analog-to-Digital Conversion Handbook, Third Edition**, Prentice-Hall, 1986.

15. W.R. Bennett, "Spectra of Quantized Signals", **Bell System Technical Journal**, No. 27, July 1948, pp. 446-472.

16. G.A. Gray and G.W. Zeoli, "Quantization and Saturation Noise Due to Analog-Digital Conversion", **IEEE Transactions on Aerospace and Electronic Systems**, Jan. 1971, pp. 222-223.

17. M.J. Tant, **The White Noise Book**, Marconi Instruments, July 1974.

18. W.A. Kester, "PCM Signal Codecs for Video Applications", **SMPTE Journal**, No. 88, November 1979, pp. 770-778.

19. Lawrence Rabiner and Bernard Gold, **Theory and Application of Digital Signal Processing**, Prentice-Hall, 1975.

20. Matthew Mahoney, **DSP -Based Testing of Analog and Mixed-Signal Circuits,** , IEEE Computer Society Press, Washington, D.C., 1987.

21. IEEE Trial-Use Standard for Digitizing Waveform Recorders, No. 1057-1988.

22. Richard J. Higgins, **Digital Signal Processing in VSLI**, Prentice-Hall, 1990.

23. **High Speed Design Seminar**, Analog Devices, 1990.

24. M. S. Ghausi and K. R. Laker, **Modern Filter Design: Active RC and Switched Capacitors**, Prentice Hall, 1981.

25. Brodersen, Gray, and Hodges, *MOS Switched-Capacitor Filters,* **Proc. IEEE Vol. 67,** January 1979, pp. 61 - 65.

26. J. T. Caves, et. al., *Sampled Analog Filtering Using Switched Capacitors as Resistor Equivalents*, **IEEE J. Solid State Circuits Vol. SC-12**, pp. 592 - 599.